

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comment regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE 01/31/01	3. REPORT TYPE AND DATES COVERED Final, 06/01/96 -- 12/31/00		
4. TITLE AND SUBTITLE Scalable multicast networks for high-performance computing and communications		5. FUNDING NUMBERS DAAH04-96-1-0234		
6. AUTHOR(S) Jianke Yang / Yuanyuan Yang		2001 MAY 16 AM 10:17		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Vermont Burlington, VT 05401				
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211		8. PERFORMING ORGANIZATION REPORT NUMBER		
		10. SPONSORING / MONITORING AGENCY REPORT NUMBER ✓ 36090-MA-DPS •48		
11. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12 b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The objective of this research is to design a class of scalable interconnection networks to support arbitrary multicast communications in highly parallel computing systems. In particular, the project was concentrated on (1) designing routing strategies for multicast networks and developing a network simulator to simulate multicast networks under routing control strategies; (2) deriving necessary and sufficient conditions under which multicast networks are nonblocking; (3) establishing analytical models for the performance of multicast networks; (4) studying self-routing scheme for multicast communication; (5) studying efficient multicast communication in optical networks, and determining rearrangeable and nonblocking conditions in optical multicast networks. (6) studying efficient multicast communication (as well as general collective communication) in both electronic and optical networks. Over the finding period, a number of important results were obtained and these results were reported in 22 journal papers, 1 book chapter, 6 patents (granted or pending), and 30 conference papers. The most important results of this project are summarized in the final progress report and all publications are also listed (attached).				
14. SUBJECT TERMS			15. NUMBER OF PAGES 10	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OR REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Final Progress Report

Project Title: Scalable Multicast Networks for High-performance Computing and Communications

ARO Award No.: DAAH04-96-1-0234

Principle Investigator: Jianke Yang/Yuanyuan Yang

Institution: University of Vermont/SUNY Stony Brook

The objective of this research is to design a class of scalable interconnection networks to support arbitrary multicast communications in highly parallel computing systems. In particular, the project was concentrated on

- Designing routing control strategies for multicast networks.
- Developing a network simulator to simulate multicast networks under routing control strategies. (2) Deriving necessary and sufficient conditions under which multicast networks are nonblocking.
- Establishing analytical models for the performance of multicast networks.
- Designing low-cost multicast networks.
- Studying efficient multicast communication in optical networks, and determining rearrangeable and nonblocking conditions in optical multicast networks.
- Establishing analytical models for the performance of fault-tolerant networks.
- Studying efficient multicast communication (as well as general collective communication) in both electronic and optical networks,
- Studying self-routing scheme for multicast communication.
- Analyzing rearrangeable and nonblocking conditions in optical multicast switching networks with multicast-capable optical switches.

The most important results of this project are summarized as follows.

- Seven routing control strategies for multicast networks were designed and compared. A multicast network simulator was developed. All seven routing control strategies were implemented on the simulator and evaluated with respect to different measures such as network utilization and blocking probability. Extensive simulations were carried out on the multicast networks for different network size,

20010607 072

workload and traffic distribution, and a large amount of statistical data were collected. The experimental finding is summarized as follows: (1) A network with comparable cost to a permutation network is almost nonblocking for multicast communication. (2) Routing control strategies are effective for reducing blocking probability of multicast networks. The best routing control strategy can provide a factor of 2 to 3 performance improvement over random routing.

- Several analytical models for the performance metrics of multicast networks, such as blocking probability, network throughput, packet transfer time and packet loss probability, were developed and compared. Two types of new multicast networks, restricted multicast networks and self-routing multicast networks, were designed. Both types of multicast networks designed have the same network cost as a permutation network of that type. The first type of network can realize a substantial amount of well-defined multicast connections in a nonblocking fashion. The second type of network can realize arbitrary multicast connections in a rearrangeable, self-routing fashion.
- Efficient multicast communication was studied for a class of WDM optical networks with direct interconnects such as linear arrays, rings, meshes, tori and hypercubes. For each type of the network, necessary and sufficient conditions for the network to be rearrangeably nonblocking and wide-sense nonblocking for multicast connections under several routing algorithms were obtained. An analytical model for the blocking probability of a fault-tolerant Clos network was developed and compared with simulation results, and our results indicate that the Clos network has good fault-tolerant capability.
- Optimal algorithms for multicast and general collective communication were proposed for two types of networks: direct networks and multistage networks in both electronic domain and optical domain. Self-routing multicast networks were designed and self-routing tag encoding schemes were studied.
- Rearrangeable and nonblocking conditions in optical multicast switching networks with multicast-capable optical switches were obtained and new designs for such type of the multicast networks were proposed and compared. A related type collective communication pattern, all-to-all communication was studied. Efficient all-to-all communication algorithms were developed for both direct networks and multistage networks.

Over the funding period, the following research papers/patents have been produced (listed in chronological order).

Refereed Journal Papers

1. Y. Yang and J. Wang, "Optimal all-to-all personalized exchange in a class of optical multistage networks," accepted for publication in *IEEE Transactions on Parallel and Distributed Systems*.
2. J. Wu and Y. Yang, "The postal network: a recursive network for parameterized communication model," accepted for publication in *Journal of Supercomputing*.
3. P. Wan, L. Liu and Y. Yang, "Optimal routing based on super topology in WDM optical parallel interconnects," accepted for publication in *Journal of Parallel and Distributed Computing*, Special Issue on Routing in Computer and Communication Networks.
4. Y. Yang, J. Wang and C. Qiao, "Nonblocking WDM multicast switching networks," *IEEE Transactions on Parallel and Distributed Systems*, vol. 11, no. 12, pp. 1274-1287, December 2000.
5. Y. Yang and J. Wang, "A more accurate analytical model on blocking probability of multicast networks," *IEEE Transactions on Communications*, vol. 48, no. 11, pp. 1930-1936, November 2000.
6. Y. Yang, "The performance of multicast banyan networks," *Journal of Parallel and Distributed Computing*, vol. 60, no. 8, pp. 909-923, August 2000.
7. P. Wan and Y. Yang, "Load balanced routing in counter-rotated SONET rings," *Networks*, vol. 35(4), pp. 279-286, 2000.
8. Y. Yang and J. Wang, "Optimal all-to-all personalized exchange in self-routable multistage networks," *IEEE Transactions on Parallel and Distributed Systems*, vol. 11, no. 3, pp. 261-274, March 2000.
9. Y. Yang, J. Wang and Y. Pan, "Permutation capability of optical multistage interconnection networks," *Journal of Parallel and Distributed Computing*, vol. 60, no. 1, pp. 72-91, Jan. 2000.
10. Y. Yang and J. Wang, "A new self-routing multicast network," *IEEE Transactions on Parallel and Distributed Systems*, vol. 10, no. 12, pp. 1299-1316, December 1999.
11. Y. Yang and G.M. Masson, "The necessary conditions for Clos-type nonblocking multicast networks," *IEEE Transactions on Computers*, vol. 48, no. 11, pp. 1214-1227, November 1999.

12. M.P. Haynos and Y. Yang, "An analytical model on the blocking probability of a fault-tolerant network," *IEEE Transactions on Parallel and Distributed Systems*, special issue on Fault-Tolerant Routing, vol 10, no. 10, pp. 1040-1051, October 1999.
13. Y. Yang, "An analytical model for the performance of buffered multicast banyan networks," *Computer Communications*, Elsevier Science Publishers, vol. 22, no. 7, pp. 598-607, May 1999.
14. Y. Yang and J. Wang, "Wide-sense nonblocking Clos networks under packing strategy," *IEEE Transactions on Computers*, vol. 48, no. 3, pp. 265-284, March 1999.
15. Y. Pan, C. Qiao and Y. Yang, "Optical multistage interconnection networks: new challenges and approaches," *IEEE Communications*, special issue on Optical Networks and Communication Systems, vol. 37, no. 2, pp. 50-56, Feb. 1999.
16. Y. Yang and N.H. Kessler, "Modeling the blocking behavior of Clos networks," *International Journal of Parallel and Distributed Systems and Networks*, vol. 2, no. 1, 1999, pp. 1-9.
17. Y. Yang, "A class of interconnection networks for multicasting," *IEEE Transactions on Computers*, vol. 47, no. 8, August 1998, pp. 899-906.
18. Y. Yang and G.M. Masson, "A fast network controller for nonblocking multicast networks," *International Journal of Parallel and Distributed Systems and Networks*, vol. 1, no. 3, 1998, pp. 149-156.
19. Y. Yang and J. Wang, "On blocking probability of multicast networks," *IEEE Transactions on Communications*, vol. 46, no. 7, July 1998, pp. 957-968.
20. L.A. Calitri and Y. Yang, "Efficient resource placement in generalized de Bruijn networks," *International Journal of Parallel and Distributed Systems and Networks*, vol. 1, no. 3, 1998, pp. 117-126.
21. Y. Yang and L.A. Calitri, "Resource placement in a class of hierarchical networks," *International Journal of Computers and Applications*, vol. 20, no. 1, 1998, pp. 15-25.
22. Y. Yang "An analytical model on network blocking probability," *IEEE Communications Letters*, vol. 1, no. 5, September 1997, pp. 143-145.

Book Chapters

23. Y. Yang, "Supporting Multicast Communication in Clos-Type Switching Networks," in *Switching Networks: Recent Advances*, D.-Z. Du and H.Q. Ngo Eds., Kluwer Academic Publishers, 2000.

Patents

24. Y. Yang and G.M. Masson, "Controller for a Non-Blocking Broadcast Network," United States Patent Number 5,801,641, issued September 1998.
25. Y. Yang and J. Wang, "Self-Router Multicast Network Architecture," United States Patent Application Serial No. 09/049,101, filed March 1998.
26. Y. Yang and J. Wang, "All-to-All Communication Scheme for Multistage Interconnection Networks," United States Patent Application filed in March 1999.
27. Y. Yang and J. Wang, "Method and Apparatus for Constructing a Latin Square Matrix for Network Communications," United States Patent Application filed in March 1999.
28. Y. Yang and J. Wang, "A Generic Approach to Generating Permutations for All-to-All Personalized Exchange for Self-Routing Multistage Interconnection Networks," United States Patent Application filed in March 1999.
29. Y. Yang and J. Wang, "Optimal All-to-All Personalized Exchange in Optical Multistage Networks," United States Patent Application filed in September 1999.

Refereed Conference Papers

30. Y. Yang and J. Wang, "Near-optimal all-to-all broadcast in multidimensional all-port meshes and tori," *Proceedings of 15th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2001)*, San Francisco, CA, April 2001.
31. Y. Wang, Y. Yang and P. Wan, "A general multicast scheme for a class of multicast-capable WDM networks," *Proceedings of 12th International Conference on Parallel and Distributed Computing and Systems (PDCS'00)*, pp. 753-760, Las Vegas, Nevada, November 2000.
32. Y. Wang and Y. Yang, "Multicasting in a class of multicast-capable WDM networks," *Proceedings of the 9th IEEE International Conference on Computer Communications and Networks (IC3N '00)*, pp. 184-191, Las Vegas, Nevada, October 2000.

33. Y. Yang, J. Wang and C. Qiao, "Nonblocking WDM multicast switching networks," *Proceedings of 2000 International Conference on Parallel Processing, (ICPP'00)*, pp. 521-530, Toronto, Canada, August 2000. IEEE Computer Society Press.
34. Y. Yang and J. Wang, "Optimal all-to-all personalized exchange in multistage networks," *Proceedings of 7th International Conference on Parallel and Distributed Systems (ICPADS'2000)*, pp. 229-236, Iwate, Japan, July 2000. IEEE Computer Society Press (Best Paper Award).
35. Y. Yang and J. Wang, "Optimal all-to-all personalized exchange in a class of optical multistage networks," *Proceedings of 14th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2000)*, pp. 753-760, Cancun, Mexico, May 2000.
36. Y. Yang and J. Wang, "All-to-all personalized exchange in banyan networks," *Proceedings of Eleventh International Conference on Parallel and Distributed Computing and Systems (PDCS'99)*, Cambridge, MA, November 1999, pp. 78-86.
37. P. Wan, L. Liu and Y. Yang, "Optimal routing based on super topology in hypercube WDM networks," *Proceedings of 24th IEEE Conference on Local Computer Networks (LCN'99)*, Lowell, MA, October 1999, pp. 142-149. IEEE Computer Society Press.
38. P. Wan and Y. Yang, "Load balancing in counter-rotated directed SONET rings," *Proceedings of 1999 International Conference on Parallel Processing (ICPP '99)*, Aizu-Wakamatsu, Japan, September 1999, pp. 542-550. IEEE Computer Society Press.
39. Y. Yang and J. Wang, "A more accurate analytical model on blocking probability of multicast networks," *Proceedings of ISCA 12th International Conference on Parallel and Distributed Computing Systems (PDCS'99)*, Fort Lauderdale, FL, August 1999, pp. 191-197.
40. Y. Yang and J. Wang, "Routing tag encoding in self-routing multicast networks," *Proceedings of 1999 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA '99)*, Las Vegas, Nevada, June 1999, pp. 2371-2377.
41. Y. Yang and J. Wang, "Efficient all-to-all broadcast in all-port mesh and torus networks," *Proceedings of 5th IEEE International Symposium on High-Performance Computer Architecture (HPCA-5)*, Orlando, FL, Jan. 1999, pp. 290-299.

42. Y. Yang (invited), "Nonblocking and almost nonblocking multicast switching networks," *Robust Communication Networks: Interconnection and Survivability*, NSF DIMACS Workshop, Rutgers University, NJ, November 1998, pp. 107-126.
43. M.P. Haynos and Y. Yang, "An analytical model on the blocking probability of a fault-tolerant network," *Proceedings of 10th International Conference on Parallel and Distributed Computing and Systems (PDCS '98)*, Las Vegas, Nevada, October 1998, pp. 363-370.
44. C. Zhou and Y. Yang, "Multicast communication in a class of wide-sense nonblocking optical WDM networks," *Proceedings of the 7th IEEE International Conference on Computer Communications and Networks (IC3N '98)*, Lafayette, Louisiana, October 1998, pp. 321-328. IEEE Computer Society Press.
45. J. Wu and Y. Yang, "The postal network: A versatile interconnection topology," *Proceedings of 1998 International Conference on Parallel Processing (ICPP '98)*, Minneapolis, MN, August 1998, pp. 612-619. IEEE Computer Society Press.
46. C. Zhou and Y. Yang, "Embedding multicast in a class of optical WDM networks," *Proceedings of 1998 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA '98)*, Las Vegas, Nevada, July 1998, pp. 597-604.
47. C. Zhou and Y. Yang, "Multicast communication in a class of rearrangeable optical WDM networks," *Proceedings of the IEEE 5th International Conference on Massively Parallel Processing Using Optical Interconnections (MPPOI '98)*, Las Vegas, Nevada, June 1998, pp. 68-75. IEEE Computer Society Press.
48. Y. Yang and J. Wang, "A new self-routing multicast network," *Proceedings of the First Merged IEEE International Parallel Processing Symposium & Symposium on Parallel and Distributed Processing (IPPS/SPDP '98)*, Orlando, FL, March 1998, pp. 351-357.
49. Y. Yang, J. Wang and Y. Pan, "Permutation capability of optical multistage interconnection networks," *Proceedings of the First Merged IEEE International Parallel Processing Symposium & Symposium on Parallel and Distributed Processing (IPPS/SPDP '98)*, Orlando, FL, March 1998, pp. 125-133.
50. C. Zhou and Y. Yang (invited), "On the number of wavelengths required to embed multicast assignments in WDM networks," *Multichannel Optical Networks: Theory and Practice*, NSF DIMACS Workshop, Rutgers University, NJ, March 1998, pp. 85-107.

51. Y. Yang, "A probabilistic model on network blocking behavior," *Proceedings of the 9th International Conference on Parallel and Distributed Computing and Systems (PDCS '97)*, Washington, D.C., October 1997, pp. 139-143.
52. C.M. Crawford and Y. Yang, "Experimental study of multicast routing strategies," *Proceedings of the 9th International Conference on Parallel and Distributed Computing and Systems (PDCS '97)*, Washington, D.C., October 1997, pp. 495-500.
53. Y. Yang, "The performance of multicast banyan networks," *Proceedings of the 6th IEEE International Conference on Computer Communications and Networks (IC3N '97)*, Las Vegas, Nevada, September 1997, pp. 51-59. IEEE Computer Society Press.
54. Y. Yang and J. Wang, "Performance analysis and simulation of multicast networks," *Proceedings of 1997 International Conference on Parallel Processing (ICPP '97)*, Bloomington, IL, August 1997, pp. 385-389. IEEE Computer Society Press.
55. Y. Yang and N.H. Kessler (invited), "Modeling the blocking behavior of Clos networks," *Advances in Switching Networks: NSF DIMACS Workshop*, Princeton, NJ, July 1997, pp. 85-102.
56. Y. Yang and J. Wang, "Wide-sense nonblocking Clos networks under packing strategy," *Proceedings of the 11th IEEE International Parallel Processing Symposium (IPPS '97)*, Geneva, Switzerland, April 1997, pp. 45-54.
57. L.A. Calitri and Y. Yang, "Resource placement in de Bruijn networks," *Proceedings of the 8th International Conference on Parallel and Distributed Computing and Systems (PDCS '96)*, Chicago, IL, October 1996, pp. 109-116.
58. Y. Yang, "A class of interconnection networks for multicasting," *Proceedings of the 10th IEEE International Parallel Processing Symposium (IPPS '96)*, Honolulu, HI, April 1996, pp. 796-802.
59. Y. Yang and G.M. Masson, "The necessary conditions for Clos-type nonblocking multicast networks," *Proceedings of the 10th IEEE International Parallel Processing Symposium (IPPS '96)*, Honolulu, HI, April 1996, pp. 789-795.

Over the funding period, a number of students have been involved in the research under my supervision and obtained advanced degrees.

• Graduate Students Supervised

- Yang Wang, Master's Degree in Computer Science, May 2000.
 Thesis Title: *Multicasting in a Class of Multicast-Capable WDM Networks*.
 Employment: i2 Technologies, Cambridge, MA.
- Matthew P. Haynos, Master's Degree in Computer Science, May 1998.
 Thesis Title: *An Analytical Model on the Blocking Probability of a Fault-Tolerant Network*.
 Employment: IBM Corporation, Encinitas, CA.
- Chunling Zhou, Master's Degree in Computer Science, Jan. 1998.
 Thesis Title: *Multicast Communication in a Class of WDM Optical Networks*.
 Employment: Sanchez Computer Associates, Inc., Malvern, PA.
- Neil H. Kessler, Master's Degree in Computer Science, May 1997.
 Thesis Title: *Analytical Model for Packing Strategy on Network Blocking Probability*.
 Employment: Thomson Financial Services, Boston, MA.
- Christopher M. Crawford, Master's Degree in Computer Science, March 1997.
 Thesis Title: *Experimental Study of Multicast Routing Strategies*.
 Employment: Sanders, Lockheed Martin, Nashua, NH.
- Lisa A. Calitri, Master's Degree in Computer Science, May 1996.
 Thesis Title: *Resource Placement in Cube-Connected-Cycles and de Bruijn Networks*.
 Employment: Clovis Point, Rochester, VT.
- Michael Landry, Master's Degree, May 1996.
 Comprehensive Exam Title: *Review of Nexus Distributed Operating System*.
 Employment: Lockheed Martin, South Burlington, VT.